Information Exploitation and Planning for a Sensor Web

Mark Abramson, Francis Carr, David Carter, Stephan Kolitz, Blair Leake, Natasha Markuzon, Peter Scheidler
The Charles Stark Draper Laboratory, Inc.
Cambridge, MA 02139

Abstract:

We present progress made in developing capabilities that fit within the EPOS sensor web dynamic replanning architecture described in our NSTC 2007 paper. We first describe a sensor web use case that is one of the drivers in our development. We then discuss two major information exploitation capabilities, the first focused on cloud cover forecasts and the second focused on wild fire prediction. We also describe several planning capabilities, including system-of-system coordination, planning for multiple heterogeneous UAVs and manned aircraft, and enhanced single satellite planning capabilities, focused on EO-1. We describe the evolution of these capabilities toward compliance with OGC standards.

I. EPOS[1] Overview

The fundamental EPOS concept of operation is that of optimized dynamic replanning and execution. Sensor data and model forecasts are inputs to a closed-loop decision-making system. In collaboration with users, EPOS monitors the input, and when appropriate, replans and executes a new plan that optimizes the tasking of available sensing assets to gather data. EPOS provides the science community with innovative capabilities that can be used to advance science modeling of the phenomena of interest. Its capabilities can also be used to provide governmental agencies and commercial interests early warning of possible hazardous situations. The high-level functional architecture for EPOS is illustrated in Figure 1.

Situation Awareness: Situation Awareness provides estimates of current world and system states.

Situation Assessment: Situation Assessment takes Situation Awareness output and uses **Information Exploitation** technologies, e.g., pattern recognition and data mining, to support monitoring, diagnosis and prediction of world and system states.

Planning and Execution: EPOS planning is optimization-based technology that will support a full range of operational autonomy, from manual operator control to full autonomy.

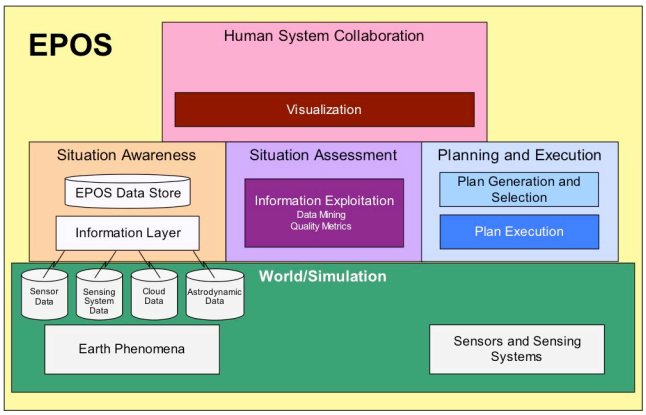

Figure 1: EPOS Functional Architecture

II. Cloud Cover Web Services Concept of Operations

We are in the early stages in the re-architecture of EPOS's situation awareness/assessment and planning capabilities as web services. At present only cloud cover estimation and forecast capability is being implemented as a web service. However, the new architecture provides a roadmap for how these capabilities will play in a web services concept of operation over time.

Figure 2 shows the concept of operations for the cloud cover web service currently under development. A user specifies a 4-D volume of interest consisting of a set of (latitude, longitude, altitude, time) 4-tuples. The user can then request values of predicted cloud cover ("forecasts") by altitude layers, and/or values of estimated actual cloud cover (i.e., cloud cover values estimated from data gathered during the time interval of interest, the "actuals"). The actuals request can be accommodated at any time after the posting time for the data from the end of the time interval of interest. Historical actuals will be available immediately. If there is no available forecast yet, i.e., the time interval of interest is more than 84 hours in the future, the response will include the times when forecasts will become available and possibly a prediction based on relative frequencies calculated from historical actuals data.

The output of the OGC[2]-compliant, EPOS cloud service will be one or more layers of cloud cover data, depending upon the altitude range specified in the 4-D volume. The

[1] Draper Lab's Earth Phenomena Observing System

[2] Open Geospatial Consortium

actuals are based on AFWA[3] WWMCA[4] data, and the forecasts on AFWA SCFM[5] data. We will also provide a quality metric and cloud probabilities associated with the forecasts. The quality metric and cloud probability calculations are the result of Draper information exploitation models and algorithms developed from analysis of the multiple terabytes of historical cloud cover data stored in the EPOS Cloud Server at Draper. Calculation of the quality metric was described in the Final Report from Draper's AIST-02 project. The calculation and evaluation of the cloud probability calculations are described in III.B.

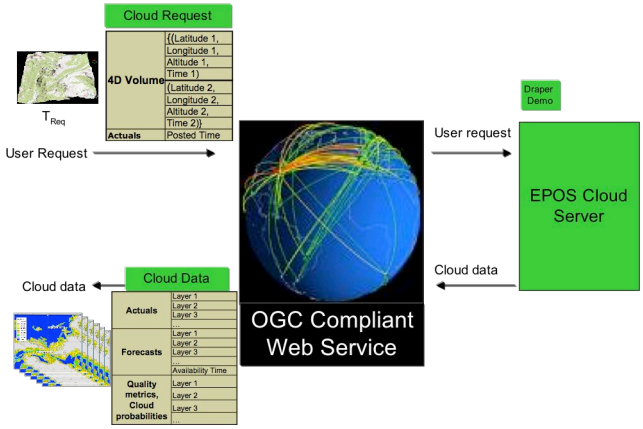

Figure 2: Cloud Cover Web Service

Results from the use of cloud forecasts in EO-1 operations are summarized in Table 1. The cumulative success rate for EPOS picks is 81.7%.

Evaluation period	January 2 - September 30, 2006 (271 days)	October 1, 2006 to March 29, 2007 (177 days)	March 30 to August 26, 2007 (153 days)	Aug 27 to April 20 2008 (237 days)
Total number of orbital revolutions with a scheduling scene	3750	2577	2117	3280
Total number of opportunities for target picks	396	179	220	375
Total number of alternate targets picked by EPOS	61	49	66	117
Total numbers of our picks that EO-1 actually imaged	55	43	61	94
Number of successful picks (less cloud cover over EPOS pick)	47 (out of 61)	37 (out of 49)	51 (out of 66)	105 (out of 117)
Number of unsuccessful picks	14 (out of 61)	12 (out of 49)	15 (out of 66)	12 (out of 117)

Table 1: Results from Ongoing EO-1 Operations using EPOS

III. Information Exploitation

There are two major Information Exploitation efforts reported in this paper, Wildfire Prediction and SCFM Analysis.

A. *Wildfire Prediction*

We have continued the development of models that predict if a wildfire will become or stay "large" the next day, given the three previous days worth of observations. We built our prediction models using data collected from the MODIS (on Terra) Thermal Anomalies/Fire Daily Level 3 sinusoidal grid product fused with NOAA daily weather data and land cover data collected from the National Land Cover Database (2001). The land cover data also allows us to estimate proximity of a fire to populated areas, providing an additional measure of fire danger.

The objective of wildfire prediction is to develop predictive modeling approaches that will improve prediction of Earth phenomena, (e.g., wildfire, harmful algae bloom, hurricanes, tornadoes). The initial application in our project is wildfire prediction.

The goal of wildfire prediction is to predict which current fires will develop into large and/or threatening ones based on several days of MODIS[6] fire observations, along with NOAA[7] weather data and Landsat land cover data.

A wildfire sensor web is a good example for demonstration of our information exploitation technology because it involves a disparate set of asynchronous distributed sensor systems with different lead times for tasking, some of which might be days in advance. For example, the target location for an image to be gathered by ASTER[8] has to be known several days in advance, and a UAV[9] used to image wildfires must file its initial flight plan at least 72 hours in advance, with constrained updates at 24 hour intervals.

We developed a model to predict if a fire on day D will become "*large*" (at least 5 contiguous fire pixels) on day D+1, and day D+2. A fire is identified based on MODIS Level 3 data – Thermal Anomalies/Fire Daily Global with 1 km resolution. Pixel values are given from 0 - 9. We considered pixels with values 7, 8, and 9 as *fire pixels*.

We fused the MODIS data with NOAA weather data consisting of daily observations of temperature, relative humidity, wind speed, wind direction, wind gusts collected from weather stations. We used land cover data from the National Land Cover Database 2001 (NLCD 2001) as a proxy for the amount of fuel available for the fire, and fused this data with the MODIS and NOAA data.

The NLCD are provided on a state-by-state basis. Figure 3 illustrates land cover of the state of California.

[3] Air Force Weather Agency
[4] World-Wide Merged Cloud Analysis
[5] Stochastic Cloud Forecast Model

[6] Moderate Resolution Imaging Spectroradiometer
[7] National Oceanic and Atmospheric Administration
[8] Advanced Spaceborne Thermal Emission and Reflection Radiometer
[9] Unmanned Air Vehicle

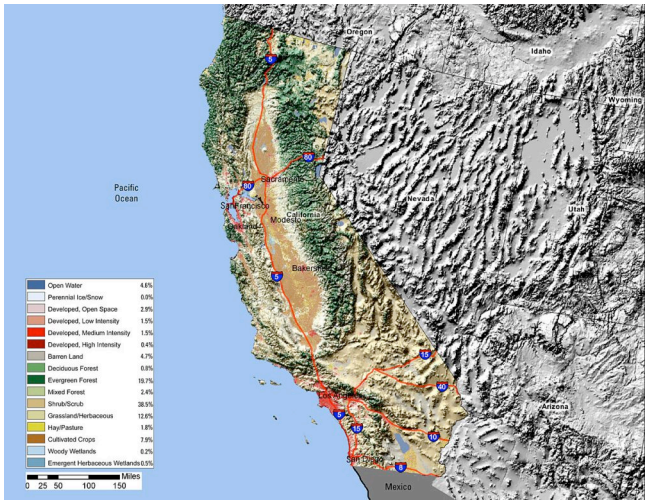

Figure 3: California Land Cover

We used historical data from 2000-2006 for the states of Arizona, California, Nevada, Utah, and Texas to develop the EPOS Wildfire Prediction Model.

A fused input data record has the following fields:

Day -2 fire info and weather
Day -1 fire info and weather
Day 0 ("today") fire info and weather
Land cover.

A representation of the data over time is given in Figure 4. Data records are color coded over the five days and show the evolution of fires. A single record is defined by processing the data from "Today." Each record in the Today file is either a large or small fire. On other days, a record could be either a large fire, small fire or no fire.

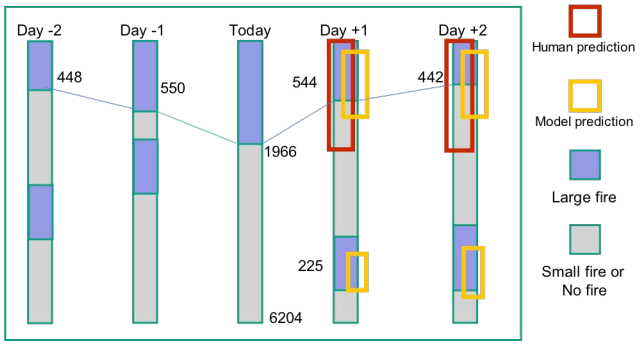

Figure 4: Fire Data for Prediction

Results from our data mining process produced a Wildfire Prediction Model, illustrated in Table 2. We found that a Random Forest Decision Tree algorithmic approach had the best performance and it is used in the Wildfire Prediction Model. It correctly predicts no large fires on Day +1 better than the human does. However, the simple use of the model in predicting large fires was equivalent to the human prediction. The human decision model is simple at this stage of our work: a large fire today will remain large for the next two days. We plan to improve the human decision model to better match reality.

			Day +1 state of the world	
			no large fire	yes large fire
Prediction	model	no large fire	80%	25%
	human		71%	26%
	model	yes large fire	20%	75%
	human		29%	74%

Table 2: Wildfire Prediction Metrics

Using the probabilistic information provided by the Wildfire Prediction Model produced results much better than human prediction. Table 3 illustrates the use of the probability of a large fire on day D+1 in prediction. If a cutoff of 0.90 or above is used, the percentage of true positives is very high and covers 39% of the cases. Because there are many more fires than sensor web resources can image, this provides excellent input into the sensor web planning function in order to use the relatively scarce sensor resources efficiently.

Probability Range	True Positives (%)	Observations (%)
Human Performance	74%	-
1.00	100%	9
0.90-0.99	91	30
0.80-0.89	78	15
0.70-0.79	70	17
0.60-0.69	63	15
0.50-0.59	59	12

Table 3: Predicting a D+1 Fire for Large Fires Active Today

B. SCFM Analysis for Use in EPOS

We access cloud data automatically from the AFWA server 24 hours a day, 7 days a week. The current cloud data (WWMCA = World-Wide Merged Cloud Analysis) is received every hour, while the forecast cloud data (SCFM = Stochastic Cloud Forecast Model) is received every six hours, approximately 1.5 hours after the nominal time of the forecast. We process the data and store in the EPOS Cloud Server. Queries by visualization and planning allow access to any of the current or forecast data sets.

We use cloud forecasts to help choose between conflicting EO-1 imaging opportunities. The simplest selection rule is this:

- *Rule 1:* When given conflicting targets of equal priority, select the target whose predicted cloud cover is smallest

A more sophisticated selection rule, which maximizes the expected number of cloud-free[10] scenes, is this:

- *Rule 2:* When given conflicting targets of equal priority, select the target for which the probability P(total cloud cover ≤ 20%) is largest

In work reported previously, we have suggested how SCFM data may be used to estimate $P(C \leq x)$, where C denotes total cloud cover. Our technique makes the following assumptions:

- Total cloud cover C at latitude θ, longitude φ is a beta-distributed random variable
- SCFM-predicted total cloud S(θ, φ) is beta-distributed with the same mean and variance[11]
- Mean and variance of S(θ, φ) are slowly-changing functions of θ and φ, so may be estimated by sampling S(θ, φ) near (θ, φ)
- The parameters α and β of the beta distribution of S(θ, φ) may be recovered from sample values using either method-of-moments or maximum likelihood estimates
- $P(C \leq x) = I(x, \alpha, \beta)$, where $I(x, \alpha, \beta) = B(x, \alpha, \beta)/B(\alpha, \beta)$ is the incomplete beta function

We have estimated $P(C \leq x)$ for a set of test cases. Figure 5 illustrates sample results obtained using this software.

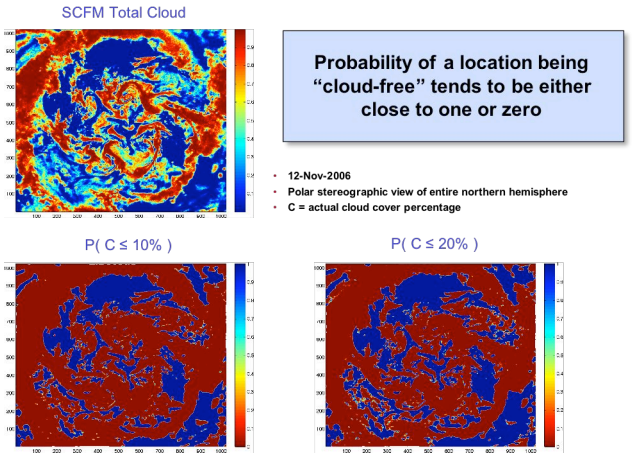

Figure 5: Sample Results using SCFM Northern Hemisphere

Note that a smaller SCFM does not imply a larger $P(C \leq 20\ \%)$. In Figure 6, the SCFM value at target location on left is 4%, while $P(C \leq 20\ \%)$ estimated from the data in the 5x5 array surrounding the target location is 54%. SCFM value at target location on right is 15%, while $P(C \leq 20\ \%)$ estimated from the data in the 5x5 array surrounding the target location is 95%. Note that a 5x5 array is used because it maps the resolution of the cloud forecast to the resolution of the underlying NOAA data used to generate the forecast.

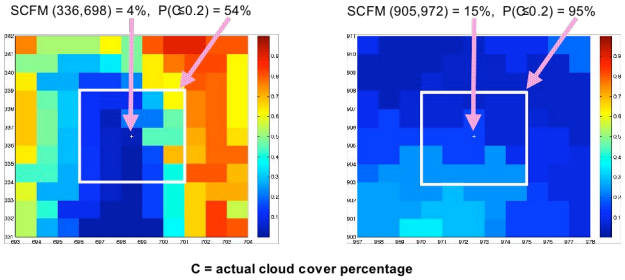

Figure 6: Smaller SCFM Does not Imply Larger P(C ≤ 20%)

We conducted an experiment to determine whether this probability estimation methodology can be used in support of EO-1 planning to increase the yield of cloud-free scenes. We used the following procedure:

- Using historical EO-1 target locations and historical EO-1 ephemeris, generate a simulated list of EO-1 imaging opportunities for a 24 hour planning period. Determine times at which targets can be imaged, limiting roll to ±2 WRS paths.
- (Rule 0) For each descending (daylight) pass, and each ascending (night) pass in the 24 hour planning period, *randomly* select one of the feasible targets. Count the number of selected targets for which WWMCA total cloud is ≤ 20%.
- (Rule 1) For each descending (daylight) pass, and each ascending (night) pass in the 24 hour planning period, select the target whose SCFM predicted total cloud is smallest. Use historical SCFM data from the forecast available 12 hours prior to the start of the 24 hour planning period. Count the number of selected targets for which WWMCA total cloud is ≤ 20%.
- (Rule 2) For each descending (daylight) pass, and each ascending (night) pass in the 24 hour planning period, select the target for which our estimate of $P(C \leq 20\ \%)$ is smallest. Use historical SCFM data from the forecast available 12 hours prior to the start of the 24 hour planning period to estimate these probabilities. Count the number of selected targets for which WWMCA total cloud is ≤ 20%.
- Repeat using historical data for different dates. Compare the performance of the two selection rules.

Here are the results of this experiment.

Number of simulated 24 hour periods: (1/1/07 to 12/30/07 at 3-day intervals)	121
Total number of day and night passes (including fractional passes):	3660
Number of Rule 0 (random) selections for which WWMCA is ≤ 20% | WWMCA is available:	1352 / 3382
Number of Rule 1 selections for which WWMCA is ≤ 20% | WWMCA is available:	1916 / 3284
Number of Rule 2 selections for which WWMCA is ≤ 20% | WWMCA is available:	2098 / 3288

[10] In this work, we call an image "cloud-free" if less than 20% of the image area is obscured by clouds.

[11] Note that SCFM-predicted total cloud is beta-distributed by construction.

Table 4 Results from Simulation Using Predicted Probability of Cloud Cover

From these results, we have the following conclusions:

- Using SCFM value to select between competing EO-1 targets (Rule 1) is significantly better than making a random selection (Rule 0) – this models current EPOS use in EO-1 operations
- Using our estimate of $P(C \leq 20\%)$ to select between competing EO-1 targets (Rule 2) yields nearly 10% more cloud-free scenes than using SCFM value alone (Rule 1) – this is what we are currently evaluating for use in EO-1 operations

We have integrated this new capability into EPOS system to support EO-1 operations. $P(C \leq 20\ \%)$ is now computed and logged for each candidate target. We plan to meet with the EO-1 mission manager and system engineer to discuss use of $P(C \leq 20\%)$ once testing is complete.

IV. PLANNING

A. PLANNING FOR MULTIPLE HETEROGENEOUS UAVS

We are extending and enhancing the single UAV planner developed last year to include multiple heterogeneous vehicles. Note that although we refer to UAVs throughout, the planning technology is equally applicable to manned aircraft.

As part of the UAV planner, a mathematical model will generate a movement and observation plan for multiple, heterogeneous UAVs. We assume that there are two types of UAVs: a larger, high-altitude airframe capable of longer distances and flight times, and a smaller airframe capable of low-altitude flights and high-resolution imaging. The model will route the UAVs on a path that will maximize the total value of the observations to the user. The UAV will observe targets that are either "point targets" or "area targets." At each target, the UAV will collect data; for example, take an aerial image of a target. An overview of the problem is shown in Figure 7. The detailed mathematical formulation is shown in Figures 8 and 9.

The inputs for the model include UAV parameters, target locations, observation value information, and operational constraints such as keep-out zones. The model will output a path for each UAV to travel and the observation times at each target.

- **General problem characteristics:**
  - All routes must be contained within a well-defined operational area, and avoid well-defined keep-out zones (population centers)
  - The output plan will be a "stick route" for each UAV.
  - Additional detailed routing may also be done as a post-processing step.
- **Model characteristics:**
  - UAVs will be assumed to have a constant speed when transiting between targets
  - UAVs will spend an amount of time at each target dependent on both the target and the UAV type

- **Model Inputs:**
  - Cruise speed
  - Image time
  - Climb/descend rate
  - Operational altitude range
  - Sensor FOV
  - Turn radius
  - Geometry of operational area and keep-out zones
- **Model Outputs:**
  - Decision variables representing which flight legs UAVs travel on (Path Plan)
  - Arrival time and duration at each target visited on plan (Observation Plan)
  - Total value to be gained from performing plan

Figure 7: UAV Planner Problem Formulation

A mathematical formulation has been developed using optimization-based models, as a mixed-integer program, as shown in Figures 8 and 9. Exact solutions may be reasonable for smaller problem instances; however, most likely the eventual solution algorithm will be a meta-heuristic; options include large-scale neighborhood search, or problem decomposition. Using a heuristic will allow the model to be solved quickly, even with large numbers of UAVs and targets. It will be able to run autonomously, with a user-interface that will allow an operator to alter the paths of the UAVs, add new targets, or even re-solve the problem with different parameters. In addition, the heuristic will be able to adjust existing solutions when new resources become available.

The solution includes a path and observation plan. The path plan will be the route for the UAVs to travel from one waypoint to the next. The observation plan will define when the UAV should arrive at each target, what activity it performs at each target, and the length of the time that the UAV will spend at the target. For example, if the UAV is observing an area target, then it could fly in a lawnmower pattern in the area.

Data:

- A, B = UAV types
  - A signifies high-alt UAV, B signifies low-alt UAV
- *a, b* = number of UAVs available of types A, B
- *I* = set of all points (targets and bases) used in routing
  - The bases will be treated as targets with no imaging time or value
- $V_i(n_A, n_B)$ = value of visiting point *i* n_A times with a UAV of type A, and n_A times with a UAV of type B.
- T_A, T_B = maximum endurance times of UAV of type A, B
- $t_{ij,A}$, $t_{ij,B}$= travel time between points *i* and *j* (including additional routing to avoid keep-out zones) required by UAVs of type A, B
- $t_{i,A}$, $t_{i,B}$ = time required to image point i, by UAVs of type A, B

Decision variables:

- $x_{ij,k}$ = 1, if the k[th] UAV of type A traverses leg from point *i* to point *j* 0 otherwise (k = 0…a)
- $y_{ij,k}$ = 1, if the k[th] UAV of type B traverses leg from point *i* to point *j* 0 otherwise (k = 0…b)

Figure 8: Model Data and Variables

$$\max z = \sum_{i \in I} V_i(n_{i,A}, n_{i,B})$$

s.t.

$$n_{i,A} = \sum_{j \in I} \sum_{k=1}^{a} x_{ji,k} \quad \forall i \in I$$

$$n_{i,B} = \sum_{j \in I} \sum_{k=1}^{b} y_{ji,k} \quad \forall i \in I$$

$$\sum_{j \in I} x_{ij,k} = \sum_{j \in I} x_{ji,k} \quad \forall i \in I, k = 1...a$$

$$\sum_{j \in I} y_{ij,k} = \sum_{j \in I} y_{ji,k} \quad \forall i \in I, k = 1...b$$

$$\sum_{i \in I} \sum_{j \in I} (t_{ij,A} + t_{j,A}) x_{ij,k} \le T_A \quad \forall k = 1...a$$

$$\sum_{i \in I} \sum_{j \in I} (t_{ij,A} + t_{j,A}) y_{ij,k} \le T_B \quad \forall k = 1...b$$

$$\sum_{j \in I} x_{ji,k} = 1 \quad j \text{ is the base}, \forall k = 1...a$$

$$\sum_{j \in I} y_{ji,k} = 1 \quad j_k \text{ is the } k^{th} \text{ base}, \forall k = 1...b$$

$$\sum_{i,j \in S} x_{ij,k} = |S| - 1 \quad \forall k = 1...a,$$

$$\forall S \subseteq I, S \neq \phi, S \text{ not including the base}$$

$$\sum_{i,j \in S} y_{ij,k} = |S| - 1 \quad \forall k = 1...b$$

$$\forall S \subseteq I, S \neq \phi, S \text{ not including the } k^{th} \text{ base}$$

$$x_{ij,k} \in \{0,1\} \quad \forall i, j \in I, \forall k \in 1...a$$

$$y_{ij,k} \in \{0,1\} \quad \forall i, j \in I, \forall k = 1...b$$

Figure 9: MIP Formulation

B. *Enhanced EO-1 Planning*

We have developed an optimization-based algorithm for enhanced EO-1 planning. In this section, we describe the algorithm in detail. We are currently writing software to test out the algorithm's performance against simpler heuristic approaches.

The planning period of interest is assumed to be 1-2 weeks in duration. Over the planning period we propagate the orbit of the EO-1 vehicle according to standard SGP4 theory to compute all of the possible viewing opportunities for each target of interest.

For each such viewing opportunity, we compute a score representing the utility that one would expect to gain if that viewing opportunity was actually attempted. Each target is described by some static information including its geographical location or boundaries, priority, and desired number of "good" images. For a given target, each viewing opportunity occurs at a distinct time and hence has its own corresponding cloud-cover forecast. For each viewing opportunity, the static target information is combined with the cloud-cover forecast to compute its score. Our algorithm does not depend on the particular formula used to compute these scores – however we do require that these scores are cumulative, i.e. the score of a *collection* of viewing opportunities is just their total score.

At this time, we are using a simple formula in which the score equals the target's priority, multiplied by the probability that the cloud-cover is less than 20%. We have also discussed a more complex formula in which the *statistical distribution* of the cloud-cover forecast is used to compute a corresponding *distribution* of utility, and the expected value of this utility is multiplied by the target's priority. An example is shown in Figure 10. We anticipate that this more complex formula will better take into account the uncertainty inherent in the forecasts.

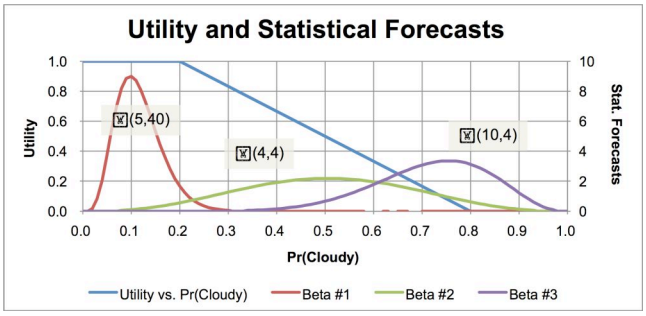

Figure 10: Utility and Statistical Forecasts

We assume that on each ascending or descending "pass" of EO-1, at most one viewing opportunity can be actually selected, which is current operational practice. While this assumption simplifies the planning problem, it could be extended to include multiple scenes on a single revolution.

After computing the scores, we visualize the planning problem as shown in the simplified diagram of Figure 11: the set of ascending and descending passes are shown on the left – the set of targets is shown on the right – each viewing opportunity is an edge (with an associated score) connecting a pass and a target. From this diagram, it is immediately apparent that the problem of selecting a subset of the possible viewing opportunities can be modeled as the well-studied *assignment problem* from operations research. Many efficient algorithms and software packages have been developed to solve the assignment problem.

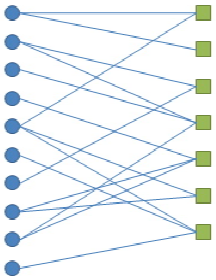

Figure 11: Formulation as an Assignment Problem

Note that the number of passes may not exactly match the number of targets, in which case extra "dummy" passes (or targets) are added as necessary, along with extra zero-score edges to connect them to all of the targets (or passes). In addition, if multiple viewings of the same target are desired, then multiple *copies* of that target are introduced, including copies of the relevant viewing opportunities.

This idea of introducing multiple copies of the same target immediately suggests a useful extension. There is uncertainty inherent in whether any given viewing opportunity will be successful. In order to gracefully plan for this uncertainty, for each target we can introduce sufficient copies so that the *expected* number of *successful* viewing opportunities is equal to the desired number of images.

C. Coordination Planner

The image request web service is how we envision the Coordination Planner we're developing will ultimately be used. This web service fits into our goal concept of how users can request images, which then initiates the processing needed to fulfill the request as well as the output provided from fulfilling the request. User requests will take the form of requirements for the image as the primary input – and could be at an abstract level such as a "theme" (e.g., wildfire, flood, hurricane), or at a more detailed level, e.g., target location, required resolution, required collection time interval, and other image requirements. The time window input will determine whether an archive image will suffice (i.e., time window start time far enough in the past) or be used in scheduling the tasking of a sensor. The priority input would be an input to the planning and scheduling of this image request. The impact of a user-specified priority will depend upon the overall priority mechanization – e.g., priorities could be a function of the theme, location, and /or the particular user..

The request will be accommodated by three coordinated services. A system-of-sensor-systems request broker service, which is functionally what we've called the Coordination Planner in previous reports, will allocate each user request to one (or more) sensor system planners. Additional detail of this service is shown in Figure 12. The allocation will be made by matching up the requirements of the user request with the capability and availability of the sensor systems under consideration. Opportunities for joint sensor system collects would be determined if simultaneous views were desired. The Coordination Planner (Figure 13) works with multiple sensor system planner services – some Draper developed, and some provided by other parties. The sensor system planner services determine what opportunities are available for meeting user requests. As discussed below, the planners generally collect multiple requests, before generating a schedule for the particular asset. The Coordination Planner will keep track of joint collection opportunities. The sensor system situational awareness and execution services accept the tasking schedules and produce feedback requested by the users. Notional feedback includes the desired image and/or situational awareness information describing what collection options are available to meet a particular user's request – from which a user can optionally add preferences.

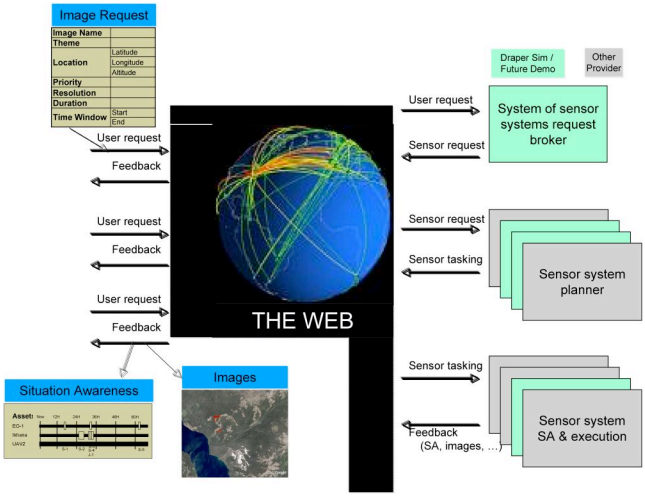

Figure 12: Image Request Web Service

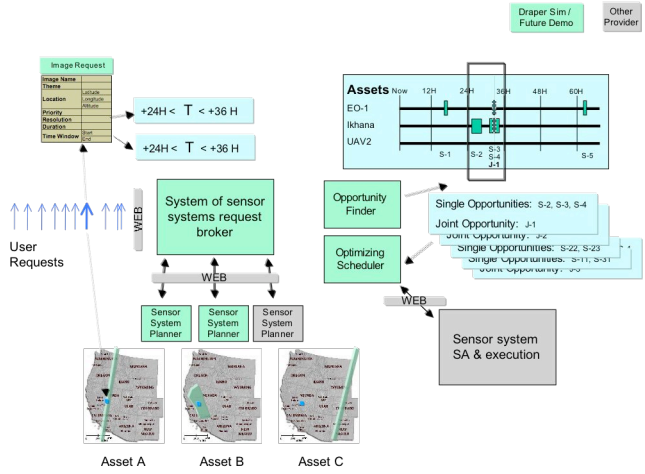

Figure 13: Coordination Planner

In general, each of the collection assets will have planning cycles that are independent of each other, as shown in Figure 14. Furthermore, the assets generally have three different periods for planning, each characterized by different rules on how new requests might get incorporated: preplanning, pre-execution planning, and execution. Note that the cycles of planning are not necessarily coordinated across assets, i.e., the ground

station planning for an overhead asset like EO-1 is not explicitly coordinated with UAV flight planning. Additionally, while an asset is getting ready to execute one plan, inputs for the next plan may be being processed (e.g., see asset C in Figure 14). The notion is that requests are being collected for an asset, and then plans may be generated at the end of a phase (e.g., Pre-planning), at fixed times within a phase, or based on events within a phase.

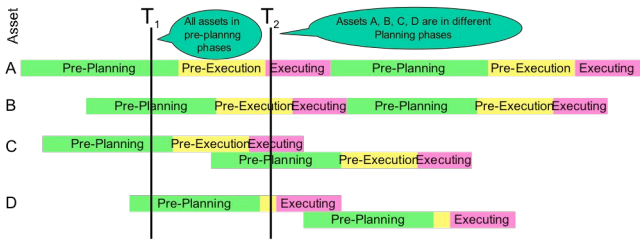

Figure 14: Planning cycles are independent across assets

Figure 15 provides further insight into how planning requests are accommodated for an individual asset through an example. The top row of arrows under the horizontal bars indicating planning phases and cycles represent a stream of user requests input to the planner. Requests are not individually addressed as they arrive at the planner. Instead, when the end of the pre-planning cycle occurs, the 10 blue arrows become requests that will be considered for cycle 1 of planning. Requests coming in after this time, shown in gray on that top row of arrows, will be accommodated by later cycles of planning. In this example, three of these requests are scheduled for execution in the first cycle – shown as pink arrows in the second row of arrows. Clouds obscure one of the targets and so only two of three actually have good images collected. Thus for the next cycle of planning, 8 requests from cycle 1 are again considered, along with 5 requests that arrived after the first planning phase was complete.

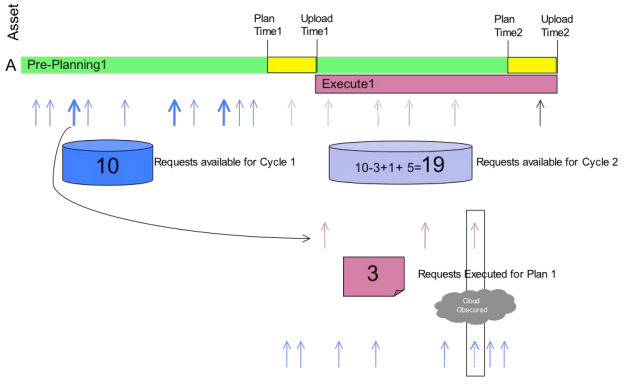

Figure 15: Planning requests over time for a single asset

Figure 16 continues this example by showing that the process exemplified in Figure 15 is carried out for all the assets making up the system.

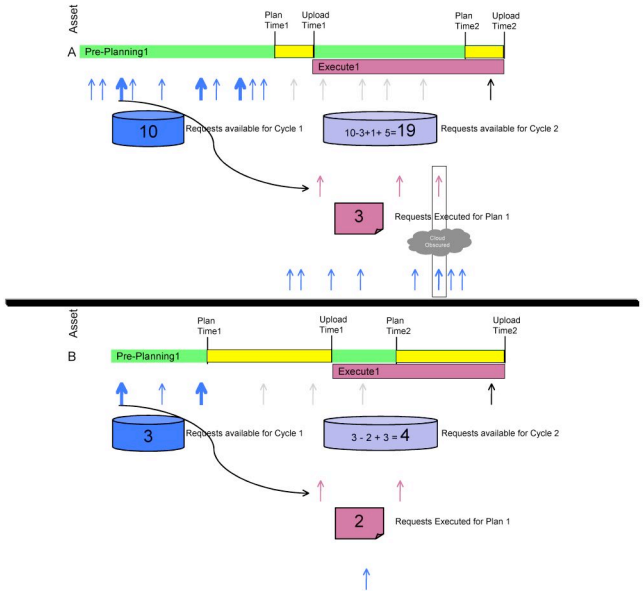

Figure 16: Planning requests over time for multiple assets

Figure 17 summarizes in one view all the web services that have been described in this subsection and how they work together to handle user requests.

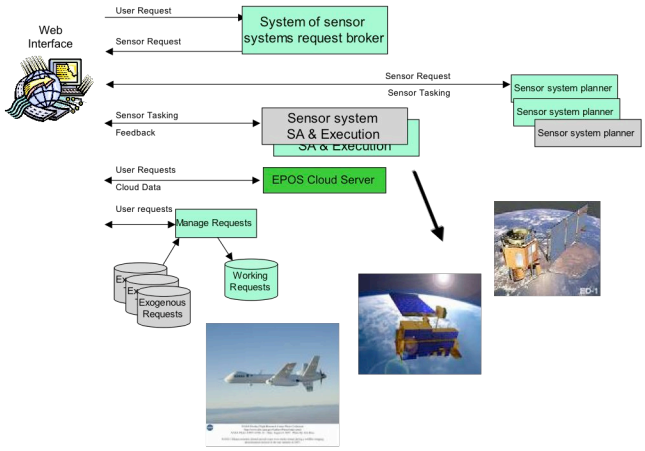

Figure 17: Architecture to support Concept of Operations

A notional view of one part of a user web service interaction is shown in Figure 18. This particular user has the ability to select which assets to be considered in planning. The more general "discovery" service that determines what assets are available for planning is not covered in this paper. The notional user screen shows the user having entered a request and having received a schedule of opportunities that could meet the request. The user is selecting one of the options to send to the sensor system SA and execution module.

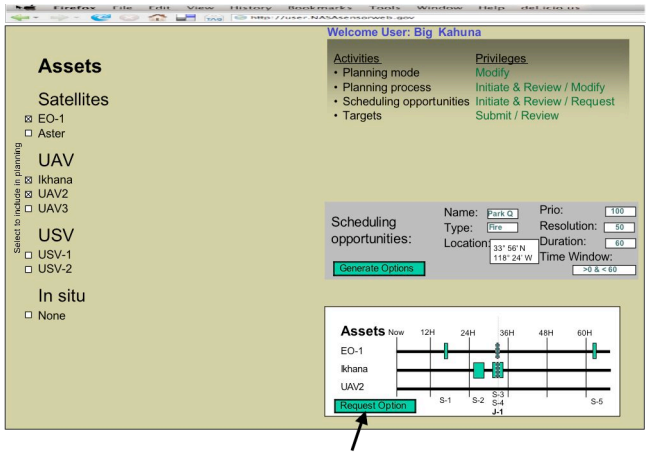

Figure 18: User Concept of Operations for Generating and Selecting Options

ACKNOWLEDGMENTS

This material is based upon work supported by the National Aeronautics and Space Administration's Earth Science Technology Office (ESTO) AIST-05 program, under Grant Number NNX06AG17G. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Aeronautics and Space Administration